

Optimal Whitening and Decorrelation

Agnan Kessy¹, Alex Lewin², and Korbinian Strimmer³ *

2 December 2015; revised 26 May 2016

¹Statistics Section, Dept. of Mathematics, Imperial College London, South Kensington Campus, London SW7 2AZ, UK.

²Dept. of Mathematics, Brunel University London, Kingstone Lane, Uxbridge UB8 3PH, UK.

³Epidemiology and Biostatistics, School of Public Health, Imperial College London, Norfolk Place, London W2 1PG, UK.

*To whom correspondence should be addressed. Email: k.strimmer@imperial.ac.uk

Abstract

Whitening, or sphering, is a common preprocessing step in statistical analysis to transform random variables to orthogonality. However, due to rotational freedom there are infinitely many possible whitening procedures. Consequently, there is a diverse range of sphering methods in use, for example based on principal component analysis (PCA), Cholesky matrix decomposition and zero-phase component analysis (ZCA), among others.

Here we provide an overview of the underlying theory and discuss five natural whitening procedures. Subsequently, we demonstrate that investigating the cross-covariance and the cross-correlation matrix between sphered and original variables allows to break the rotational invariance and to identify optimal whitening transformations. As a result we recommend two particular approaches: ZCA-cor whitening to produce sphered variables that are maximally similar to the original variables, and PCA-cor whitening to obtain sphered variables that maximally compress the original variables.

Keywords: Whitening, decorrelation, ZCA-Mahalanobis transformation, principal components analysis, Cholesky decomposition, CAT score, CAR score.

1 Introduction

Whitening, or *sphering*, is a linear transformation that converts a d -dimensional random vector $\mathbf{x} = (x_1, \dots, x_d)^T$ with mean $E(\mathbf{x}) = \boldsymbol{\mu} = (\mu_1, \dots, \mu_d)^T$ and positive definite $d \times d$ covariance matrix $\text{var}(\mathbf{x}) = \boldsymbol{\Sigma}$ into a new random vector

$$\mathbf{z} = (z_1, \dots, z_d)^T = \mathbf{W}\mathbf{x} \quad (1)$$

of the same dimension d and with unit diagonal “white” covariance $\text{var}(\mathbf{z}) = \mathbf{I}$. The square $d \times d$ matrix \mathbf{W} is called the whitening matrix. As orthogonality among random variables greatly simplifies multivariate data analysis both from a computational and a statistical standpoint, whitening is a critically important tool, most often employed in preprocessing but also as part of modeling (Zuber and Strimmer, 2009; Hao et al., 2015).

Whitening can be viewed as a generalization of *standardizing* a random variable which is carried out by

$$\mathbf{z} = \mathbf{V}^{-1/2}\mathbf{x}, \quad (2)$$

where the matrix $\mathbf{V} = \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$ contains the variances $\text{var}(x_i) = \sigma_i^2$. This results in $\text{var}(z_i) = 1$ but it does not remove correlations. Often, standardization and whitening transformations are also accompanied by mean-centering of \mathbf{x} or \mathbf{z} to ensure $E(\mathbf{z}) = 0$, but this is not actually necessary for producing unit variances or a white covariance.

The whitening transformation Equation (1) requires the choice of a suitable whitening matrix \mathbf{W} . From $\text{var}(\mathbf{z}) = \mathbf{W}\boldsymbol{\Sigma}\mathbf{W}^T = \mathbf{I}$ it follows that the matrix \mathbf{W} needs to satisfy the condition

$$\mathbf{W}^T\mathbf{W} = \boldsymbol{\Sigma}^{-1}. \quad (3)$$

However, this constraint does not uniquely determine the whitening matrix. Quite the contrary, given $\boldsymbol{\Sigma}$ there still remain infinitely many possible matrices \mathbf{W} and corresponding whitening transformations, all of which produce orthogonal but different whitened random variables.

This raises two important issues: first, how to best understand the differences among the various sphering transformations, and second, how to select an optimal whitening procedure for a particular situation. Here, we propose to address these questions by investigating the cross-covariance and cross-correlation matrix between \mathbf{z} and \mathbf{x} . As a result, we identify five natural whitening procedures, of which we generally recommend two particular approaches.

2 Notation and useful identities

In the following we will make use of a number of covariance matrix identities: the decomposition $\boldsymbol{\Sigma} = \mathbf{V}^{1/2}\mathbf{P}\mathbf{V}^{1/2}$ of the covariance matrix into the correlation matrix \mathbf{P} and the diagonal variance matrix \mathbf{V} , and the eigen-decomposition of the covariance matrix $\boldsymbol{\Sigma} = \mathbf{U}\boldsymbol{\Lambda}\mathbf{U}^T$ and the eigen-decomposition of the correlation matrix $\mathbf{P} = \mathbf{G}\boldsymbol{\Theta}\mathbf{G}^T$, where \mathbf{U} , \mathbf{G} contain the eigenvectors and $\boldsymbol{\Lambda}$, $\boldsymbol{\Theta}$ the eigenvalues of $\boldsymbol{\Sigma}$, \mathbf{P} respectively.

Following the standard convention we assume that the eigenvalues are sorted in order from largest to smallest value. In addition, we will frequently use $\Sigma^{-1/2} = \mathbf{U}\Lambda^{-1/2}\mathbf{U}^T$, the unique inverse matrix square root of Σ , as well as $\mathbf{P}^{-1/2} = \mathbf{G}\Theta^{-1/2}\mathbf{G}^T$, the unique inverse matrix square root of the correlation matrix.

3 Rotational freedom in whitening

The constraint Equation (3) on the whitening matrix does not fully identify \mathbf{W} but allows for rotational freedom. This becomes apparent by writing \mathbf{W} in its polar decomposition

$$\mathbf{W} = \mathbf{Q}_1 \Sigma^{-1/2}, \quad (4)$$

where \mathbf{Q}_1 is an orthogonal matrix with $\mathbf{Q}_1^T \mathbf{Q}_1 = \mathbf{I}_d$. Clearly, \mathbf{W} satisfies Equation (3) regardless of the choice of \mathbf{Q}_1 .

This implies a geometrical interpretation of whitening as a combination of multivariate rescaling by $\Sigma^{-1/2}$ and rotation by \mathbf{Q}_1 . It also shows that all whitening matrices \mathbf{W} have the same singular values $\Lambda^{-1/2}$, which follows from the singular value decomposition $\mathbf{W} = (\mathbf{Q}_1 \mathbf{U}) \Lambda^{-1/2} \mathbf{U}^T$ with $\mathbf{Q}_1 \mathbf{U}$ orthogonal. This highlights that the fundamental rescaling is via the square root of the eigenvalues $\Lambda^{-1/2}$. Geometrically, the whitening transformation with $\mathbf{W} = \mathbf{Q}_1 \mathbf{U} \Lambda^{-1/2} \mathbf{U}^T$ is a rotation \mathbf{U}^T followed by scaling, possibly followed by another rotation (depending on choice of \mathbf{Q}_1).

Since in many situations it is desirable to work with standardized variables $\mathbf{V}^{-1/2} \mathbf{x}$ another useful decomposition of \mathbf{W} that also directly demonstrates the inherent rotational freedom is

$$\mathbf{W} = \mathbf{Q}_2 \mathbf{P}^{-1/2} \mathbf{V}^{-1/2}, \quad (5)$$

where \mathbf{Q}_2 is a further orthogonal matrix with $\mathbf{Q}_2^T \mathbf{Q}_2 = \mathbf{I}_d$. Evidently, this \mathbf{W} also satisfies the constraint of Equation (3) regardless of the choice of \mathbf{Q}_2 .

In this view, with $\mathbf{W} = \mathbf{Q}_2 \mathbf{G} \Theta^{-1/2} \mathbf{G}^T \mathbf{V}^{-1/2}$, the variables are first scaled by the square root of the diagonal variance matrix, then rotated by \mathbf{G}^T , then scaled again by the square root of the eigenvalues of the correlation matrix, and possibly rotated once more (depending on the choice of \mathbf{Q}_2).

For the above two representations to result in the same whitening matrix \mathbf{W} two different rotations \mathbf{Q}_1 and \mathbf{Q}_2 are required. These are linked by $\mathbf{Q}_1 = \mathbf{Q}_2 \mathbf{A}$ where the matrix $\mathbf{A} = \mathbf{P}^{-1/2} \mathbf{V}^{-1/2} \Sigma^{1/2}$ is itself orthogonal. Since the eigen-decompositions of the covariance and the correlation matrix are not readily related to each other, the matrix \mathbf{A} can unfortunately not be further simplified.

4 Cross-covariance and cross-correlation

For studying the properties of the different whitening procedures we will now focus on two particularly useful quantities, namely the *cross-covariance* and *cross-correlation* matrix between the whitened vector \mathbf{z} and the original random vector \mathbf{x} . As it turns

out, these are closely linked to the rotation matrices Q_1 and Q_2 encountered in the above two decompositions of W (Equation (4) and Equation (5)).

The cross-covariance matrix Φ between z and x is given by

$$\begin{aligned}\Phi &= \text{cov}(z, x) = \text{cov}(Wx, x) \\ &= W\Sigma = Q_1\Sigma^{1/2}.\end{aligned}\tag{6}$$

Likewise, the cross-correlation matrix is

$$\begin{aligned}\Psi &= \text{cor}(z, x) = \Phi V^{-1/2} \\ &= Q_2A\Sigma^{1/2}V^{-1/2} = Q_2P^{1/2}.\end{aligned}\tag{7}$$

Thus, we find that the rotational freedom inherent in W , which is represented by the matrices Q_1 and Q_2 , is directly reflected in the corresponding cross-covariance Φ and cross-correlation Ψ between z and x . This provides the leverage that we will use to select and discriminate among whitening transformations by appropriately choosing or constraining Φ or Ψ .

Furthermore, as can be seen from Equation (6) and Equation (7), both Φ and Ψ are in general not symmetric, unless $Q_1 = I$ or $Q_2 = I$, and the diagonal elements of the cross-correlation matrix Ψ need not be equal to 1.

5 Five natural whitening procedures

In practical application of whitening there are a handful of sphering procedures that are most commonly used (e.g. Li and Zhang, 1998). Accordingly, in Table (1) we describe the properties of five whitening transformations, listing the respective sphering matrix W , the associated rotation matrices Q_1 and Q_2 , and the resulting cross-covariances Φ and cross-correlations Ψ . All five methods are natural whitening procedures arising from specific constraints on Φ or Ψ , as we will show further below.

The *ZCA whitening* transformation employs the sphering matrix

$$W^{\text{ZCA}} = \Sigma^{-1/2}.\tag{8}$$

Table 1: Five natural whitening transformations and their properties.

	Sphering matrix W	Cross- covariance Φ	Cross- correlation Ψ	Rotation matrix Q_1	Rotation matrix Q_2
ZCA	$\Sigma^{-1/2}$	$\Sigma^{1/2}$	$\Sigma^{1/2}V^{-1/2}$	I	A^T
PCA	$\Lambda^{-1/2}U^T$	$\Lambda^{1/2}U^T$	$\Lambda^{1/2}U^TV^{-1/2}$	U^T	U^TA^T
Cholesky	L^T	$L^T\Sigma$	$L^T\Sigma V^{-1/2}$	$L^T\Sigma^{1/2}$	$L^TV^{1/2}P^{1/2}$
ZCA-cor	$P^{-1/2}V^{-1/2}$	$P^{1/2}V^{1/2}$	$P^{1/2}$	A	I
PCA-cor	$\Theta^{-1/2}G^TV^{-1/2}$	$\Theta^{1/2}G^TV^{1/2}$	$\Theta^{1/2}G^T$	G^TA	G^T

where ZCA stands for “zero-phase components analysis” (Bell and Sejnowski, 1997). This procedure is also known as *Mahalanobis whitening*. With $Q_1 = I$ it is the unique sphering method with a symmetric whitening matrix.

PCA whitening is based on scaled principal component analysis (PCA) and uses

$$W^{\text{PCA}} = \Lambda^{-1/2} U^T \quad (9)$$

(e.g. Friedman, 1987). This transformation first rotates the variables using the eigenmatrix of the covariance Σ as is done in standard PCA. This results in orthogonal components, but with in general different variances. To achieve whitened data the rotated variables are then scaled by the square root of the eigenvalues $\Lambda^{-1/2}$. PCA whitening is probably the most widely applied whitening procedure due to its connection with PCA.

It can be seen that the PCA and ZCA whitening transformations are related by a rotation U , so ZCA whitening can be interpreted as rotation followed by scaling followed by the rotation U back to the original coordinate system. The ZCA and the PCA sphering methods both naturally follow the polar decomposition of Equation (4), with Q_1 equal to I and U^T respectively.

Another widely known procedure is *Cholesky whitening* which is based on Cholesky factorization of the precision matrix $LL^T = \Sigma^{-1}$. This leads to the sphering matrix

$$W^{\text{Chol}} = L^T \quad (10)$$

where L is the unique lower triangular matrix with positive diagonal values. The same matrix L can also be obtained from a QR decomposition of $W^{\text{ZCA}} = (\Sigma^{1/2} L) L^T$.

A further approach is the *ZCA-cor whitening* transformation, which is for example used in the correlation-adjusted variable selection methods involving the CAT and CAR test statistics (Zuber and Strimmer, 2009; Ahdesmäki and Strimmer, 2010; Zuber and Strimmer, 2011; Zuber et al., 2012). This approach employs

$$W^{\text{ZCA-cor}} = P^{-1/2} V^{-1/2} \quad (11)$$

as its sphering matrix. The ZCA-cor whitening method arises from first standardizing the random variable by multiplication with $V^{-1/2}$ and subsequently employing ZCA whitening based on the correlation rather than covariance matrix. The resulting whitening matrix $W^{\text{ZCA-cor}}$ differs from W^{ZCA} , and unlike the latter it is in general asymmetric.

In a similar fashion, *PCA-cor whitening* is conducted by applying PCA whitening to standardized variables. This approach uses

$$W^{\text{PCA-cor}} = \Theta^{-1/2} G^T V^{-1/2} \quad (12)$$

as its sphering matrix. Here, the standardized variables are rotated by the eigenmatrix of the correlation matrix, followed by scaling using the correlation eigenvalues. Note that $W^{\text{PCA-cor}}$ differs from W^{PCA} .

PCA-cor whitening has the same relation to the ZCA-cor transformation as does PCA whitening to the ZCA transformation. Specifically, ZCA-cor whitening can be

interpreted as PCA-cor whitening followed by a rotation G back to the frame of the standardized variables. Both the ZCA-cor and the PCA-cor transformation naturally follow the decomposition of Equation (5), with Q_2 equal to I and G^T respectively.

Finally, we may also apply the Cholesky whitening transformation to standardized variables. However, this does not lead to a new whitening procedure, as the resulting sphering matrix remains identical to W^{Chol} since the Cholesky factor of the inverse correlation matrix P^{-1} is $V^{1/2}L$, and therefore $W^{\text{Chol-cor}} = (V^{1/2}L)^T V^{-1/2} = L^T = W^{\text{Chol}}$.

6 Optimal whitening

We now demonstrate how an optimal sphering matrix W , and hence an optimal whitening approach, can be identified by evaluating suitable objective functions computed from the cross-covariance Φ and cross-correlation Ψ . Intriguingly, for each of the five natural whitening transforms listed in Table (1) we find a corresponding optimality criterion.

6.1 ZCA-Mahalanobis whitening

In many applications of whitening it is desirable to remove correlations with minimal additional adjustment, with the aim that the transformed variable z remains as similar as possible to the original vector x .

One possible implementation of this idea is to find the whitening transformation that minimizes the total squared distance between the original and whitened variables (e.g. Eldar and Oppenheim, 2003). Using mean-centered random vectors z_c and x_c with $E(z_c) = 0$ and $E(x_c) = 0$ this least squares objective can be expressed as

$$\begin{aligned} E \left((z_c - x_c)^T (z_c - x_c) \right) &= \text{tr}(I) - 2E \left(\text{tr} \left(z_c x_c^T \right) \right) + \text{tr}(\Sigma) \\ &= d - 2\text{tr}(\Phi) + \text{tr}(V). \end{aligned} \quad (13)$$

Since the dimension d and sum of the variances $\text{tr}(V) = \sum_{i=1}^d \sigma_i^2$ do not depend on the whitening matrix W minimizing Equation (13) is equivalent to maximizing the trace of the cross-covariance matrix

$$\text{tr}(\Phi) = \sum_{i=1}^d \text{cov}(z_i, x_i) = \text{tr} \left(Q_1 \Sigma^{1/2} \right) \equiv g_1(Q_1). \quad (14)$$

Proposition 1. *Maximization of $g_1(Q_1)$ uniquely determines the optimal whitening matrix to be the symmetric sphering matrix W^{ZCA} .*

Proof. $g_1(Q_1) = \text{tr}(Q_1 U \Lambda^{1/2} U^T) = \text{tr}(\Lambda^{1/2} U^T Q_1 U) \equiv \text{tr}(\Lambda^{1/2} B) = \sum_i \Lambda_{ii}^{1/2} B_{ii}$ since Λ is diagonal. As Q_1 and U are both orthogonal $B \equiv U^T Q_1 U$ is also orthogonal. This implies diagonal entries $B_{ii} \leq 1$, with equality signs for all i occurring only if $B = I$,

hence the maximum of $g_1(\mathbf{Q}_1)$ is assumed at $\mathbf{B} = \mathbf{I}$, or equivalently at $\mathbf{Q}_1 = \mathbf{I}$. From Equation (4) it follows that the corresponding optimal sphering matrix is $\mathbf{W} = \mathbf{\Sigma}^{-1/2} = \mathbf{W}^{\text{ZCA}}$. \square

For related proofs see also Genizi (1993, p. 412) and Garthwaite et al. (2012, p. 789).

As a result, we find that *ZCA-Mahalanobis whitening* is the unique procedure that *maximizes the average cross-covariance* between each component of the whitened and original vectors. Furthermore, with $\mathbf{Q}_1 = \mathbf{I}$ it is also the unique whitening procedure with a symmetric cross-covariance matrix $\mathbf{\Phi}$.

6.2 ZCA-cor whitening

In the optimization using Equation (13) the underlying similarity measure is the cross-covariance between the whitened and original random variables. This results in an optimality criterion that depends on the variances and hence on the scale of the original variables. An alternative scale-invariant objective can be constructed by comparing the centered whitened variable with the centered *standardized* vector $\mathbf{V}^{-1/2}\mathbf{x}_c$. This leads to the minimization of

$$\mathbb{E} \left((\mathbf{z}_c - \mathbf{V}^{-1/2}\mathbf{x}_c)^T (\mathbf{z}_c - \mathbf{V}^{-1/2}\mathbf{x}_c) \right) = 2d - 2\text{tr}(\mathbf{\Psi}). \quad (15)$$

Equivalently, we can maximize instead the trace of the cross-correlation matrix

$$\text{tr}(\mathbf{\Psi}) = \sum_i^d \text{cor}(z_i, x_i) = \text{tr} \left(\mathbf{Q}_2 \mathbf{P}^{1/2} \right) \equiv g_2(\mathbf{Q}_2). \quad (16)$$

Proposition 2. *Maximization of $g_2(\mathbf{Q}_2)$ uniquely determines the whitening matrix to be the asymmetric sphering matrix $\mathbf{W}^{\text{ZCA-cor}}$.*

Proof. Completely analogous to Proposition 1, we can write $g_2(\mathbf{Q}_2) = \text{tr}(\mathbf{Q}_2 \mathbf{G} \mathbf{\Theta}^{1/2} \mathbf{G}^T) = \sum_i \mathbf{\Theta}_{ii}^{1/2} C_{ii}$ where $\mathbf{C} \equiv \mathbf{G}^T \mathbf{Q}_2 \mathbf{G}$ is orthogonal. By the same argument as before it follows that $\mathbf{Q}_2 = \mathbf{I}$ maximizes $g_2(\mathbf{Q}_2)$. From Equation (5) it follows that $\mathbf{W} = \mathbf{P}^{-1/2} \mathbf{V}^{-1/2} = \mathbf{W}^{\text{ZCA-cor}}$. \square

As a result, we identify *ZCA-cor whitening* as the unique procedure that *ensures that the components of the whitened vector \mathbf{z} remain maximally correlated with the corresponding components of the original variables \mathbf{x}* . In addition, with $\mathbf{Q}_2 = \mathbf{I}$ it is also the unique whitening transformation exhibiting a symmetric cross-correlation matrix $\mathbf{\Psi}$.

6.3 PCA whitening

Another frequent aim in whitening is the generation of new uncorrelated variables \mathbf{z} that are useful for dimension reduction and data compression. In other words, we would like to construct components z_1, \dots, z_d such that the first few components in \mathbf{z} represent as much as possible the variation present in the all original variables x_1, \dots, x_d .

One way to formalize this is to use the sum of squared cross-covariances $\phi_i = \sum_{j=1}^d \phi_{ij}^2 = \sum_{j=1}^d \text{cov}(z_i, x_j)^2$ between each individual z_i and all x_j as a measure of how effectively each z_i integrates, or compresses, the original variables. Note that here, unlike in ZCA-Mahalanobis whitening, the objective function links each component in \mathbf{z} simultaneously with all components in \mathbf{x} . In vector notation the ϕ_i can be more elegantly written as

$$(\phi_1, \dots, \phi_d)^T = \text{diag}(\Phi \Phi^T) = \text{diag}(Q_1 \Sigma Q_1^T) \equiv \mathbf{h}_1(Q_1). \quad (17)$$

Our aim is to find a whitened vector \mathbf{z} such that the ϕ_i are maximized with $\phi_i \geq \phi_{i+1}$.

Proposition 3. *Maximization of $\mathbf{h}_1(Q_1)$ subject to monotonically decreasing ϕ_i is achieved by the whitening matrix \mathbf{W}^{PCA} .*

Proof. The vector $\mathbf{h}_1(Q_1)$ can be written as $\text{diag}(Q_1 \Sigma Q_1^T) = \text{diag}(Q_1 \mathbf{U} \Lambda \mathbf{U}^T Q_1^T)$. Setting $Q_1 = \mathbf{U}^T$ we arrive at $\mathbf{h}_1(\mathbf{U}^T) = \text{diag}(\Lambda)$, i.e. for this choice the ϕ_i corresponding to each component of z_i are equal to the corresponding eigenvalues of Σ . As the eigenvalues are already sorted in decreasing order, we find (cf. Table (1)) that whitening with \mathbf{W}^{PCA} leads to a sphered variable \mathbf{z} with monotonically decreasing ϕ_i . For general Q_1 the i -th element of $\mathbf{h}_1(Q_1)$ is $\sum_j \Lambda_{jj} D_{ij}^2$ where $\mathbf{D} \equiv Q_1 \mathbf{U}$ is orthogonal. This is maximized when $\mathbf{D} = \mathbf{I}$, or equivalently, $Q_1 = \mathbf{U}^T$. \square

As a result, *PCA whitening* is singled out as the unique sphering procedure that *maximizes the integration, or compression, of all components of the original vector \mathbf{x} in each component of the sphered vector \mathbf{z} based on the cross-covariance Φ as underlying measure*. Thus, the fundamental property of PCA that principal components are optimally ordered with respect to dimension reduction (Jolliffe, 2002) carries over also to PCA whitening.

6.4 PCA-cor whitening

For reasons of scale-invariance we prefer to optimize cross-correlations rather than cross-covariances for whitening with compression in mind. This leads to the sum of squared cross-correlation $\psi_i = \sum_{j=1}^d \psi_{ij}^2 = \sum_{j=1}^d \text{cor}(z_i, x_j)^2$ as measure of integration and compression, and correspondingly to the objective function

$$(\psi_1, \dots, \psi_d)^T = \text{diag}(\Psi \Psi^T) = \text{diag}(Q_2 \mathbf{P} Q_2^T) = \mathbf{h}_2(Q_2). \quad (18)$$

Proposition 4. *Maximization of $\mathbf{h}_2(Q_2)$ subject to monotonically decreasing ψ_i is achieved by using $\mathbf{W}^{\text{PCA-cor}}$ as the sphering matrix.*

Proof. Analogous to the proof of Proposition 3 we find $Q_2 = \mathbf{G}^T$ to yield optimal and decreasing ψ_i and with Equation (5) we arrive at $\mathbf{W} = \Theta^{-1/2} \mathbf{G}^T \mathbf{V}^{-1/2} = \mathbf{W}^{\text{PCA-cor}}$. \square

Hence, the *PCA-cor whitening* transformation is the unique transformation that *maximizes the integration, or compression, of all components of the original vector \mathbf{x} in each component of the sphered vector \mathbf{z} employing the cross-correlation Ψ as underlying measure*.

6.5 Cholesky whitening

Finally, we investigate the connection between Cholesky whitening and corresponding characteristics of the cross-covariance and cross-correlation matrices. Unlike the other four whitening methods listed in Table (1), which result from optimization, Cholesky whitening is due to a symmetry constraint.

Specifically, the whitening matrix \mathbf{W}^{Chol} leads to a cross-covariance matrix Φ that is lower-triangular with positive diagonal elements as well as to a cross-correlation matrix Ψ with the same properties. This is a consequence of the Cholesky factorization with \mathbf{L} being subject to the same constraint. Crucially, as \mathbf{L} is unique the converse argument is valid as well, and hence *Cholesky whitening* is the unique whitening procedure that results from lower-triangular positive diagonal cross-covariance and cross-correlation matrices.

A consequence of using Cholesky factorization for whitening is that we implicitly assume an ordering of the variables. This can be useful specifically in time course analysis to account for auto-correlation (cf. Pourahmadi, 2011, and references therein).

7 Application

7.1 Data-based whitening

In the sections above, we have discussed the theoretical background of whitening in terms of random variables x and z and using the population covariance Σ to guide the construction of a suitable sphering matrix \mathbf{W} .

In practice, however, we frequently need to whiten data rather than random variables. In this case we have an $n \times d$ data matrix $\mathbf{X} = (x_{ki})$ whose rows are assumed to be drawn from a distribution with expectation μ and covariance matrix Σ . In this setting the transformation of Equation (1) from original to whitened data matrix becomes $\mathbf{Z} = \mathbf{X}\mathbf{W}^T$.

A further complication is that the covariance matrix Σ is often unknown. Accordingly, it needs to be learned from data, either from \mathbf{X} or from another suitable data set, yielding a covariance matrix estimate $\hat{\Sigma}$. Typically, for large sample size n and small dimension d the standard unbiased empirical covariance $\mathbf{S} = (s_{ij})$ with $s_{ij} = \frac{1}{n-1} \sum_{k=1}^n (x_{ki} - \bar{x}_i)(x_{kj} - \bar{x}_j)$ is used. In high-dimensional cases with $p > n$ the empirical estimator breaks down, and the covariance matrix needs to be estimated by a suitable regularized method instead (e.g. Schäfer and Strimmer, 2005; Pourahmadi, 2011). Finally, from the spectral decomposition of the estimated covariance or corresponding correlation matrix we then obtain the desired estimated whitening matrix $\hat{\mathbf{W}}$.

7.2 Iris flower data example

For an illustrative comparison of the five natural whitening transforms discussed in this paper and listed in Table (1) we applied them on the well-known iris flower data set of Anderson reported in Fisher (1936), which comprises $d = 4$ correlated variables

Table 2: Whitening transforms applied to the iris flower data set.

	ZCA	PCA	Cholesky	ZCA-cor	PCA-cor
$\widehat{\text{cor}}(z_1, x_1)$	0.7137	0.8974	0.3760	0.8082	0.8902
$\widehat{\text{cor}}(z_2, x_2)$	0.9018	0.8252	0.8871	0.9640	-0.8827
$\widehat{\text{cor}}(z_3, x_3)$	0.8843	-0.0121	0.2700	0.6763	0.0544
$\widehat{\text{cor}}(z_4, x_4)$	0.5743	0.1526	1.0000	0.7429	0.0754
$\text{tr}(\widehat{\Phi})$	2.9829	1.1978	1.9368	2.8495	0.5059
$\text{tr}(\widehat{\Psi})$	3.0742	1.8632	2.5331	3.1914	0.1373
$\max \text{diag}(\widehat{\Phi}\widehat{\Phi}^T)$	3.1163	4.2282	3.9544	1.7437	4.1885
$\max \text{diag}(\widehat{\Psi}\widehat{\Psi}^T)$	1.9817	2.8943	2.7302	1.0000	2.9185

Bold font indicates best method, and italic font the second best method.

(x_1 : sepal length, x_2 : sepal width, x_3 : petal length, x_4 : petal width) and $n = 150$ observations.

The results are shown in Table (2) with all estimates based on the empirical covariance \mathbf{S} . The upper half shows the estimated cross-correlations between each component of the whitened and original vector for the five methods, and the lower half the values of the various objective functions discussed above.

As expected, the ZCA and the ZCA-cor whitening produce sphered variables that are most correlated to the original data on a component-wise level, with the former achieving the best fit for the covariance-based and the latter for the correlation-based objective.

In contrast, the PCA and PCA-cor methods are best at producing whitened variables that are maximally simultaneously linked with all components of the original variables. Consequently, as can be seen from the top half of Table (2), for PCA and PCA-cor whitening only the first two components z_1 and z_2 are highly correlated with their respective counterparts x_1 and x_2 , whereas the subsequent pairs z_3, x_3 and z_4, x_4 are effectively uncorrelated. Furthermore, the last line of Table (2) shows that PCA-cor whitening achieves higher maximum total squared correlation of the first component z_1 with all components of x than PCA whitening, indicating better compression.

Interestingly, Cholesky whitening always assumes third place in the rankings, either behind ZCA and ZCA-cor whitening, or behind PCA and PCA-cor whitening. Moreover, it is the only approach where by construction one pair (z_4, x_4) perfectly correlates between whitened and original data.

8 Whitening-based variable ranking and selection

The optimality properties of ZCA-cor whitening directly lead to a deeper understanding of the variable ranking procedure using CAT and CAR scores described in Zuber and Strimmer (2009, 2011), also discussed in Allen and Tibshirani (2012), Section 4.

We have d original predictors x_1, \dots, x_d , a response y , and whitened predictors z_1, \dots, z_d , and we would like to rank the covariates with regard to their importance to predict y . If all predictors are uncorrelated then variable ranking is trivial as the natural measure of variable importance is, in case of continuous y , the squared marginal correlation between each predictor and the response (Zuber and Strimmer, 2011), and for categorical y the squared t -scores between each group mean and the overall pooled mean (Ahdesmäki and Strimmer, 2010).

The main conceptual idea behind the CAT and CAR scores is to first whiten the data using the ZCA-cor transformation to remove correlation among predictors, and subsequently to conduct univariate variable ranking using the resulting orthogonal z . The CAR scores are defined as the correlations $\text{cor}(z_i, y)$ between each whitened predictor and response y , and likewise the CAT scores are the t -scores computed for the transformed predictors. Crucially, employing the specific sphering $W^{\text{ZCA-cor}}$ of Equation (11) ensures that the whitened predictors z remain maximally correlated with the original predictors x , thereby retaining interpretability when variables are selected, for example by cross-validation.

Interestingly, for computing CAT and CAR scores we do not actually need to leave the original scale. Rather than explicitly conducting ZCA-cor whitening we can, equivalently, obtain the t -scores and marginal correlations for the original correlated predictors x and subsequently “correlation-adjust” the resulting test statistics by multiplication with the factor $P^{-1/2}$ (Zuber and Strimmer, 2009, 2011).

Another strategy for variable ranking and selection based on whitening is discussed in Hao et al. (2015). In this approach, which is a variant of principal components regression (Massy, 1965), the predictors are sphered using PCA whitening followed by variable selection using lasso regression. These authors are primarily concerned with prediction and achieving sparsity, hence their choice of a compression-based whitening procedure.

9 Conclusion

In this note we have investigated linear transformations for whitening of random variables. These methods are commonly employed in data analysis for preprocessing and to facilitate subsequent analysis.

In principle, there are infinitely many possible whitening procedures all satisfying the fundamental constraint of Equation (3) for the underlying whitening matrix. However, as we have demonstrated here, the rotational freedom inherent in whitening can be broken by considering cross-covariance Φ and cross-correlations Ψ between whitened and original variables.

Specifically, we have studied five natural whitening transforms, cf. Table (1), all of which can be interpreted as either optimizing a suitable function of Φ or Ψ , or satisfying a symmetry constraint on Φ or Ψ . As a result, this not only leads to a better understanding of the differences among whitening methods, but also enables an informed choice.

In particular, selecting a suitable whitening transformation depends on the context

of application, specifically whether minimal adjustment or compression of data is desired. In the former the whitened variables remain highly correlated to the original variables, and thus maintain their original interpretation. This is advantageous for example in the context of variable selection where one would like to understand the resulting selected submodel. In contrast, in a compression context the whitened variables by construction bear no interpretable relation to the original data but instead reflect their intrinsic effective dimension.

In general, we advocate using scale-invariant optimality functions and thus recommend using cross-correlation Ψ as a basis for optimization. Consequently, we particularly endorse two specific whitening approaches. If the aim is to obtain sphered variables that are maximally similar to the original ones, we suggest to employ the ZCA-cor whitening procedure of Equation (11). Conversely, if maximal compression is desirable we recommend to use the PCA-cor whitening approach of Equation (12).

References

- Ahdesmäki, M. and Strimmer, K. (2010). Feature selection in omics prediction problems using cat scores and false non-discovery rate control. *Ann. Appl. Statist.*, 4:503–519.
- Allen, G. I. and Tibshirani, R. (2012). Inference with transposable data: modelling the effects of row and column correlations. *J. R. Statist. Soc. B*, 74:721–743.
- Bell, A. J. and Sejnowski, T. J. (1997). The “independent components” of natural scenes are edge filters. *Vision Res.*, 37:3327–3338.
- Eldar, Y. C. and Oppenheim, A. V. (2003). MMSE whitening and subspace whitening. *IEEE Trans. Inf. Theory*, 49:1846–1851.
- Fisher, R. A. (1936). The use of multiple measurements in taxonomic problems. *Ann. Eugenics*, 7:179–188.
- Friedman, J. H. (1987). Exploratory projection pursuit. *J. Am. Stat. Assoc.*, 82:249–266.
- Garthwaite, P. H., Critchley, F., Anaya-Izquierdo, K., and Mubwandarikwa, E. (2012). Orthogonalization of vectors with minimal adjustment. *Biometrika*, 99:787–798.
- Genizi, A. (1993). Decomposition of R^2 in multiple regression with correlated regressors. *Statistica Sinica*, 3:407–420.
- Hao, N., Dong, B., and Fan, J. (2015). Sparsifying the Fisher linear discriminant by rotation. *J. R. Statist. Soc. B*, 77:827–851.
- Jolliffe, I. T. (2002). *Principal Component Analysis*. Springer, New York, 2nd edition.
- Li, G. and Zhang, J. (1998). Sphering and its properties. *Sankhya A*, 60:119–133.

- Massy, W. F. (1965). Principal components regression in exploratory statistical research. *J. Am. Stat. Assoc.*, 60:234–246.
- Pourahmadi, M. (2011). Covariance estimation: the GLM and regularization perspectives. *Stat. Sci.*, 26:369–387.
- Schäfer, J. and Strimmer, K. (2005). A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics. *Stat. Appl. Genet. Molec. Biol.*, 4:32.
- Zuber, V., Duarte Silva, A. P., and Strimmer, K. (2012). A novel algorithm for simultaneous SNP selection in high-dimensional genome-wide association studies. *BMC Bioinformatics*, 13:284.
- Zuber, V. and Strimmer, K. (2009). Gene ranking and biomarker discovery under correlation. *Bioinformatics*, 25:2700–2707.
- Zuber, V. and Strimmer, K. (2011). High-dimensional regression and variable selection using CAR scores. *Stat. Appl. Genet. Molec. Biol.*, 10:34.